

WHAT IS CLAIMED IS:

1. A method of processing a database query command in a distributed database system in which a plurality of database tables are stored on a plurality of nodes each having a plurality of processors and a shared memory router, different portions of at least one database table being stored by the plurality of processors on the plurality of nodes, the method comprising:

receiving a database query command at a first node;

generating a join table for each of a plurality of processors on said first node in accordance with said database query command, said join table being generated from a portion of a database table stored by each of said plurality of processors on said first node;

sending a first message having a single copy of said join table from a first shared memory router on said first node to a second shared memory router on a second node;

storing said single copy of said join table in a common memory of said second node; and

sending a second message to a plurality of processors on said second node indicating the location of said single copy of said join table stored in said common memory.

2. The method of Claim 1, further comprising comparing said single copy of said join table stored in said common memory by each of said plurality of processors on said second node to generate a plurality of intermediate results files.

3. The method of Claim 2, further comprising sending said plurality of intermediate results files from said second shared memory router to said first shared memory router.

4. The method of Claim 3, further comprising generating a final results file from said plurality of intermediate results files.

5. The method of Claim 4, further comprising executing post-processing operations on said final results file.

6. The method of Claim 1, wherein said second message comprises a memory pointer.

7. The method of Claim 1, wherein said portions of said database table are stored by each of said plurality of processors in substantially equal portions.

8. The method of Claim 7, wherein said portion of said database table are stored by each of said plurality of processors in substantially equal portions according to a round robin distribution.

9. The method of Claim 1, wherein said storing of said portions of said database table are stored on a volatile memory of said first and second nodes.

10. The method of Claim 1, further comprising storing said portions of said database table on a persistent storage device.

11. A distributed database system for processing a database query command in which a plurality of database tables are stored on a plurality of nodes each having a plurality of processors and a shared memory router, different portions of at least one database table being stored by the plurality of processors on the plurality of nodes, the system comprising:

- a first node configured to receive a database query command;

- a plurality of processors on said first node configured to generate a join table in accordance with said database query command, said join table being generated from a portion of a database table stored by each of said plurality of processors on said first node;

- a first shared memory router on said first node configured to send a first message having a single copy of said join table; and

- a second shared memory router on a second node configured to

- receive said first message and store said single copy of said join table in a common memory of said second node, and

- send a second message to a plurality of processors on said second node indicating the location of said single copy of said join table stored in said common memory.

12. The method of Claim 11, wherein said plurality of processors on said second node are configured to compare said single copy of said join table stored in said common memory and to generate a plurality of intermediate results files.

13. The method of Claim 12, wherein said second shared memory router is further configured to send said plurality of intermediate results files to said first shared memory router.

14. The method of Claim 13, further comprising a primary controller on said first node configured to generate a final results file from said plurality of intermediate results files.

15. The method of Claim 14, wherein said primary controller is further configured to execute post-processing operations on said final results file.

16. The method of Claim 11, wherein said second message comprises a memory pointer.

17. The method of Claim 11, wherein said plurality of processors are configured to store portions of said database table in substantially equal portions.

18. The method of Claim 17, wherein said plurality of processors are configured to store portions of said database table in substantially equal portions according to a round robin distribution.

19. The method of Claim 11, wherein said plurality of processors are configured to store said portions of said database table on a volatile memory of said first and second nodes.

20. The method of Claim 11, wherein said plurality of processors are configured to store said portions of said database table on a persistent storage device.

21. A system for transmitting data between a plurality of nodes in a distributed computing system having a plurality of processors, the system comprising:

- a first node having a first shared memory router, said first shared memory router being configured to send a first data message identifying a task to be performed; and

- a second node having a second shared memory router, said second shared memory router being configured to

- receive said first data message,

- store said first data message in a common memory of said second node, and

send a second data message to at least one processor on said second node indicating the location of said first data message stored in said common memory, said second data message comprising a pointer to said first data message in said common memory of said second node.

22. The method of Claim 21, wherein said first data message comprises a memory pointer.

23. The method of Claim 22, wherein said first node is configured to store said first data message on a volatile memory of said first node.

24. The method of Claim 21, wherein said first node is configured to store said first data message on a volatile memory of said first node.